IPv6 Configuration in Linux

Nick Schmalenberger

November 23, 2006

ARPA and BBN

- Advanced Research Projects Agency of the federal Department of Defense

ARPA and BBN

- Advanced Research Projects Agency of the federal Department of Defense
- wanted a non-hierarchical computer network resistant to nuclear attack

ARPA and BBN

- Advanced Research Projects Agency of the federal Department of Defense
- wanted a non-hierarchical computer network resistant to nuclear attack
- Bolt Beranek and Newman hired to build first routers and administer network

IAB

- network originally called arpanet interconnected with others: inter-net

IAB

- network originally called arpanet interconnected with others: inter-net
- Internet Architecture Board formed in 1983

IAB

- network originally called arpanet interconnected with others: inter-net
- Internet Architecture Board formed in 1983
- responsible for RFCs and other internet design documents

IPv4

- First mentioned in 1979, in IEN documents

IPv4

- First mentioned in 1979, in IEN documents
- RFC791 + updates are used throughout 1980s and later

IPv4

- First mentioned in 1979, in IEN documents
- RFC791 + updates are used throughout 1980s and later
- Most widely used Internet Protocol

IPv4

- First mentioned in 1979, in IEN documents
- RFC791 + updates are used throughout 1980s and later
- Most widely used Internet Protocol
- 32 bit addresses - too small, ran out

IPv4

- First mentioned in 1979, in IEN documents
- RFC791 + updates are used throughout 1980s and later
- Most widely used Internet Protocol
- 32 bit addresses - too small, ran out
- Few provisions for quality-of-service (QOS) traffic prioritization or built-in encryption
- Classful routing - network determination of an address based on first two bits causes inefficient address allocation

Improvements and successors to IPv4

- QOS features available with RSVP specified in RFC2205 + updates

Improvements and successors to IPv4

- QOS features available with RSVP specified in RFC2205 + updates
- Classless InterDomain Routing (CIDR) does away with classes and allows more efficient address block assignment
- CIDR introduces slash subnet notation, replacing dotted decimal
- Variable Length Subnet Masking (VLSM) has similar benefits to CIDR within an administrative organization

Improvements and successors to IPv4

- QOS features available with RSVP specified in RFC2205 + updates
- Classless InterDomain Routing (CIDR) does away with classes and allows more efficient address block assignment
- CIDR introduces slash subnet notation, replacing dotted decimal
- Variable Length Subnet Masking (VLSM) has similar benefits to CIDR within an administrative organization
- Network Address Translation (NAT) allows the redundant use of private address space as specified in RFC1918

Improvements and successors to IPv4

- QOS features available with RSVP specified in RFC2205 + updates
- Classless InterDomain Routing (CIDR) does away with classes and allows more efficient address block assignment
- CIDR introduces slash subnet notation, replacing dotted decimal
- Variable Length Subnet Masking (VLSM) has similar benefits to CIDR within an administrative organization
- Network Address Translation (NAT) allows the redundant use of private address space as specified in RFC1918
- SSL, IPSec, and other forms of tunneling provide encrypted networking with IPv4

Improvements and successors to IPv4

- QOS features available with RSVP specified in RFC2205 + updates
- Classless InterDomain Routing (CIDR) does away with classes and allows more efficient address block assignment
- CIDR introduces slash subnet notation, replacing dotted decimal
- Variable Length Subnet Masking (VLSM) has similar benefits to CIDR within an administrative organization
- Network Address Translation (NAT) allows the redundant use of private address space as specified in RFC1918
- SSL, IPSec, and other forms of tunneling provide encrypted networking with IPv4

Improvements and successors to IPv4

- QOS features available with RSVP specified in RFC2205 + updates
- Classless InterDomain Routing (CIDR) does away with classes and allows more efficient address block assignment
- CIDR introduces slash subnet notation, replacing dotted decimal
- Variable Length Subnet Masking (VLSM) has similar benefits to CIDR within an administrative organization
- Network Address Translation (NAT) allows the redundant use of private address space as specified in RFC1918
- SSL, IPSec, and other forms of tunneling provide encrypted networking with IPv4
- proposals for IPv6, previously known as IPng for next generation, considered in 1992. Recommendation given in 1994.

IPv6 from 1994 to present

- 6bone network set up for IPv6 experiments

IPv6 from 1994 to present

- 6bone network set up for IPv6 experiments
- Special /16 subnet 3FFE designated for 6bone

IPv6 from 1994 to present

- 6bone network set up for IPv6 experiments
- Special /16 subnet 3FFE designated for 6bone
- Tunnels across the IPv4 internet are used to connect 6bone sites

IPv6 from 1994 to present

- 6bone network set up for IPv6 experiments
- Special /16 subnet 3FFE designated for 6bone
- Tunnels across the IPv4 internet are used to connect 6bone sites
- In this time most vendors develop IPv6 support in their network software

IPv6 from 1994 to present

- 6bone network set up for IPv6 experiments
- Special /16 subnet 3FFE designated for 6bone
- Tunnels across the IPv4 internet are used to connect 6bone sites
- In this time most vendors develop IPv6 support in their network software
- 6bone phased out as specified in RFC3701, production IPv6 addresses go into general use

IPv6 today

- Tunnel brokers allocating production addresses, some also offer direct links or colocation

IPv6 today

- Tunnel brokers allocating production addresses, some also offer direct links or colocation
- US federal agencies are requiring IPv6 support in new network equipment purchased

IPv6 today

- Tunnel brokers allocating production addresses, some also offer direct links or colocation
- US federal agencies are requiring IPv6 support in new network equipment purchased
- Asian and European countries adopting IPv6 faster than United States because the United States allocated more addresses to itself in the beginning of IPv4

Addresses

- Addresses are 128 bits long

Addresses

- Addresses are 128 bits long
- Hexadecimal is used with blocks of 16 bits separated by a colon. Two colons are allowed only once to indicate a portion of the address that is all zeros. Hexadecimal digits at the beginning of a block that would be zero may also be dropped, as with IPv4. Such abbreviations are to make writing long addresses easier for people, and complete addresses may be necessary in certain situations such as making a reverse DNS PTR record.

Addresses

- When putting an IPv6 address in a URL with an application port specified, brackets are used to separate the IPv6 block colons from the colon between the address and the port number.

Addresses

- When putting an IPv6 address in a URL with an application port specified, brackets are used to separate the IPv6 block colons from the colon between the address and the port number.
- An example of an IPv6 address using abbreviation in a URL is http://[2001:db8::1]:8080/
- In this address, the 2001:0DB8::/32 network is used, which is reserved for use in documentation.

Addresses

- IPv6 has the fe8::/10 network designated for use as "link-local addresses" similar to 169.254.0.0/16 in IPv4 as specified in RFC3330. These would be used with a suffix based on the link address. This address would be used in requesting the global prefix for the site from a router, which would then have the suffix appended. If no router is available, the link-local address could continue to be used in communicating with other hosts on the link.

Addresses

- IPv6 has the fe8::/10 network designated for use as "link-local addresses" similar to 169.254.0.0/16 in IPv4 as specified in RFC3330. These would be used with a suffix based on the link address. This address would be used in requesting the global prefix for the site from a router, which would then have the suffix appended. If no router is available, the link-local address could continue to be used in communicating with other hosts on the link.
- DHCP also supports IPv6 which is useful in case of netbooting

Addresses

- IPv6 has the fe8::/10 network designated for use as "link-local addresses" similar to 169.254.0.0/16 in IPv4 as specified in RFC3330. These would be used with a suffix based on the link address. This address would be used in requesting the global prefix for the site from a router, which would then have the suffix appended. If no router is available, the link-local address could continue to be used in communicating with other hosts on the link.
- DHCP also supports IPv6 which is useful in case of netbooting
- The localhost address in IPv6 is ::1

Addresses

- Backbone ISPs are assigned /32 networks. Smaller ISPs receive /48 networks from the backbone ISPs. Individuals receive /64 networks from their ISP. The /64 network address can be the global network prefix for a site advertised by a router for IPv6 hosts autoconfiguring on the network.

Addresses

- Backbone ISPs are assigned /32 networks. Smaller ISPs receive /48 networks from the backbone ISPs. Individuals receive /64 networks from their ISP. The /64 network address can be the global network prefix for a site advertised by a router for IPv6 hosts autoconfiguring on the network.
- ISPs allocated /32 may give out smaller or larger networks than /48 depending on their circumstances. Because much of their network is hierarchical and not connected to other backbones, flexibility in their address assignment also allows for more effective route summarization.

Addresses

- Backbone ISPs are assigned /32 networks. Smaller ISPs receive /48 networks from the backbone ISPs. Individuals receive /64 networks from their ISP. The /64 network address can be the global network prefix for a site advertised by a router for IPv6 hosts autoconfiguring on the network.
- ISPs allocated /32 may give out smaller or larger networks than /48 depending on their circumstances. Because much of their network is hierarchical and not connected to other backbones, flexibility in their address assignment also allows for more effective route summarization.
- The sixxs.net website offers the Ghost Route Hunter tool showing the status of various IPv6 allocations.

IPv6 Routing

- Because of the more hierarchical nature of address assignment in IPv6 as compared to IPv4, aggregate route advertisement makes routing tables in backbone routers smaller, which compensates somewhat for the increased size of the addresses.

IPv6 Routing

- Because of the more hierarchical nature of address assignment in IPv6 as compared to IPv4, aggregate route advertisement makes routing tables in backbone routers smaller, which compensates somewhat for the increased size of the addresses.
- Even though classful routing was done away with in IPv4 with CIDR, the current IPv6 allocation scheme using /32, /48, and /64 is similar to the classful system. However, because full prefix length must always be specified,

Multicast, Anycast, Unicast, and Broadcast packets

- In IPv6, many things that are done with broadcast addresses in IPv4 are done with anycast addresses.

Multicast, Anycast, Unicast, and Broadcast packets

- In IPv6, many things that are done with broadcast addresses in IPv4 are done with anycast addresses.
- Other things are done with multicast packets. To determine the link layer address of a network address, the suffix is appended to a designated multicast prefix and is sent out. The response will have the correct link address which is then cached.

Tunnels

- Two types of tunnels are mainly used

Tunnels

- Two types of tunnels are mainly used
- IPv6-in-IPv4 tunnels are used on computers with global ip addresses

Tunnels

- Two types of tunnels are mainly used
- IPv6-in-IPv4 tunnels are used on computers with global ip addresses
- UDP tunnels are used for connecting to IPv6 networks from inside a NAT private network
- UDP tunnels are necessary because if the global ip address of the NAT router changes, a tcp tunnnel would break

Linux IPv6 utilities

- Special versions of ping, traceroute, and tracepath are used

Linux IPv6 utilities

- Special versions of ping, traceroute, and tracepath are used
- iproute, route, and ifconfig all support IPv6 with options